

TELECOMMUNICATION NETWORK SYNCHRONISATION

Field of the Invention

- 5 The present invention relates to the synchronisation of nodes in a telecommunication network and in particular, though not necessarily, to the synchronisation of nodes in a Universal Mobile Telecommunications System network.

Background to the Invention

10

In a digital communication network, such as a telecommunications network or a private network having several private branch exchanges, it is often necessary to synchronise the time clocks of respective network nodes in order to ensure correct operation of the network. Network synchronisation permits all nodes on the network to operate from a common time base. This means that when one node (i.e. an intersection point) sends data to another node, both nodes can be expected to operate at approximately the same rate ensuring the successful transfer of data between the nodes. Background information on the need for network node synchronisation can be found in EP0450828.

15

- 20 In so-called "master-slave" synchronisation, one master node is chosen to distribute high quality clock signals (generated by a Primary Reference Clock (PRC)) to all slave nodes in a hierarchy of network nodes. The master node distributes PRC clock signals to adjacent nodes which in turn distribute the received and regenerated clock signals to their adjacent nodes until all the nodes in the network are using the same clock origin.

25

The need for synchronisation is especially important in mobile telecommunication networks, and will become even more so with the introduction of Universal Mobile Telecommunications System (UMTS) networks where the UMTS Terrestrial Radio Access Network (UTRAN) places very severe limits on network synchronisation.

30

A typical UTRAN configuration consists of Radio Network Controllers (RNCs) which perform switching functions in the network (analogous in some ways with conventional telephone exchanges and with Mobile Switching Centres of GSM networks) and Radio

09737324-124400

Base Stations (RBSs) which provide the interface between the UTRAN and the mobile terminals (each RBS being responsible for a given cell). The RNCs and RBSs are arranged in a hierarchy (or hierarchies) with a single RNC possibly being responsible for tens of RBSs. The link structure in a UTRAN may be complex, with nodes of the same type being linked to one another as well as to nodes of a different type. In certain circumstances, synchronisation may be taken from a co-located GSM network or UTRAN synchronisation may be utilised in GSM nodes.

In the event of a synchronisation failure, e.g. due to the failure of a link between two nodes, action must be taken quickly to re-establish synchronisation. This usually means selecting for the node suffering from the effects of the failure (as well as for other nodes downstream of that node) an alternative incoming link which can be used to achieve synchronisation. Typically, certain incoming links are preferred to other links for this purpose, and the selection of an appropriate link requires a network level administration system which is connected to all network nodes. This work requires each node of the network to have a complete knowledge of the network and, in failure situations, the network synchronisation can suffer from unforeseen combinations of the network nodes.

WO95/24801 describes a method of synchronising a network by propagating synchronisation messages down through a hierarchy of network nodes. The synchronisation messages each comprise a master node address, a distance-to-master node, indicated as the number of intermediate nodes through which the message has passed, and the identity of the transmitting node. Each node through which a message passes, increases a distance counter by 1 and changes the transmitting node identity to its own identity. The path field allows receiving nodes to prioritise incoming links for synchronisation purposes.

WO96/39760 describes a method of detecting timing loops in a Synchronous Digital Hierarchy (SDH) network by sending a synchronisation message consisting of the identities of all the nodes through which the synchronisation message has passed. The synchronisation message also contains a count of the number of nodes through which clock signal has passed. This is used to prevent excessively long synchronisation chains.

Summary of the Invention

It is an object of the present invention to overcome or at least mitigate the disadvantages of known synchronisation networks. In particular, it is an object of the present invention to provide a synchronisation network in which synchronisation problems may be overcome substantially on a node level, automatically and with no or minimal operator intervention. It is a second object of the present invention to allow for the fast stabilisation of a network synchronisation process. It is a third object of the present invention to allow newly introduced network nodes to be rapidly synchronised with the network.

According to a first aspect of the present invention there is provided a method of synchronising nodes of a telecommunication network in which a master node is coupled to a Primary Reference Clock (PRC) and a plurality of slave nodes are each arranged to synchronise their internal clock to the PRC using data received on incoming data link, the method comprising:

propagating Synchronisation Status Messages through the network from the master node, with each node through which a message passes incorporating into the message its own identity, thereby generating in each message a path which has been followed by the message;

introducing a delay in the propagation of the messages at at least certain of the network nodes; and

for each of at least some of the incoming links of each node, registering the path and/or path length of a Synchronisation Status Message received on a link as an attribute for that link.

Embodiments of the present invention allow a node to compare the merits of different incoming data links as sources of synchronisation information. In the event that synchronisation (or re-synchronisation) is required, the node may select that incoming link having an attribute indicating the shortest path length from the master node. The introduction of a delay in the propagation of messages at at least certain nodes, increases the probability that a synchronisation message will be received first at a given

node over a shorter path, rather than over a longer path. This will tend to decrease the overall time taken to synchronise the network.

It will be appreciated that it is necessary to propagate Synchronisation Status Messages on initialising a new network. Synchronisation Status Messages may also be broadcast periodically or at other intervals thereafter in order to enable the network to cope with dynamic changes in network architecture (e.g. due to the failure of an inter-node link or the introduction of a new link or node).

Synchronisation Status Messages may be generated in response to receipt at the master node of a Synchronisation Status Request Message sent from another network node. Such a Request Message may be sent by a new node upon introduction to the network. A Synchronisation Status Message may be generated by a slave node in response to receipt at that slave node of a Synchronisation Status Request Message sent from a neighbouring slave node, with the Synchronisation Status Message including an identification of the path over which the sending slave node has been synchronised.

A node through which a Synchronisation Status Message passes may additionally add to the message its own "distance" from the master node. This distance may be defined by way of the number of node-to-node hops made by the message to get from the master node to the current node. Nodes adjacent to the master node have a distance of $PRC + 1$, nodes adjacent to nodes having a distance of $PRC + 1$ have a distance of $PRC + 2$, etc. For each incoming link, a node may register the distance included in a Synchronisation Status Message received on that link as an attribute for that link.

The present invention is particularly applicable to mobile telecommunications networks such as GSM and UMTS (more particularly to the UTRAN part of a UMTS network). However, the invention is also applicable to fixed line networks such as Public Switched Telephone Networks (PSTNs).

The delay introduced by a slave node may be the same for all slave nodes which introduce a delay. Alternatively, the delay may increase with distance from the master node. Preferably, slave nodes neighbouring the master node do not introduce a delay.

The delay to be introduced by a node may be incorporated into a Synchronisation Status Message. This avoids the need to have delay tables at all network nodes. However, in the alternative, delay tables may be present at all nodes.

5 According to a second aspect of the present invention there is provided a telecommunications network comprising a master node coupled to a Primary Reference Clock (PRC) and a plurality of slave nodes, each of the slave nodes being arranged to synchronise their internal clock to the PRC using data received on incoming data link, each of the slave nodes comprising:

10 means for receiving on the or each of at least some of the incoming links to the node, a Synchronisation Status Message incorporating the identities of the nodes through which the message has passed;

means for registering the path or path length of the Synchronisation Status Message as an attribute for the link on which it was received;

15 means for incorporating into one of said messages the identity of the node, thereby generating in the message a path which has been followed by the message; and

means for propagating the modified Synchronisation Status Message to neighbouring nodes using outgoing links,

20 wherein at least certain of the nodes in the network are arranged to introduce a delay in the propagation of the respective modified messages.

It will be appreciated that a receiving node will synchronise on the best incoming link, as identified by the paths of the Synchronisation Status Messages received on the incoming links. The Synchronisation Status Message received on the best incoming
25 link is the message to which the node will incorporate its identity, and which is propagated to the neighbouring nodes.

According to a third aspect of the present invention there is provided a node for use in a multi-node telecommunications network, the node comprising:

30 means for receiving on the or each incoming link to the node a Synchronisation Status Message incorporating the identities of the nodes through which the message has passed;

09737324-121400

means for registering the path or path length of the or each Synchronisation Status Message as an attribute for the link on which it was received;

means for incorporating into one of the messages the identity of the node, thereby generating in the message a node path which has been followed by the message;

5 and

means for propagating the modified Synchronisation Status Message to neighbouring nodes using outgoing links, after a predefined time delay.

10 According to a fourth aspect of the present invention there is provided a method of synchronising nodes of a telecommunication network in which a master node is coupled to a Primary Reference Clock (PRC) and a plurality of slave nodes are each arranged to synchronise their internal clock to the PRC using data received on incoming data link, the method comprising:

15 propagating Synchronisation Status Messages through the network from the master node, with each slave node through which a message passes incrementing a distance counter contained in the message, thereby generating in each message a path length taken by the message;

introducing a delay in the propagation of the messages at at least certain of the network nodes; and

20 for each of at least some of the incoming links of each node, registering the path length of a Synchronisation Status Message received on a link as an attribute for that link.

25 According to a fifth aspect of the present invention there is provided a method of synchronising nodes of a telecommunication network in which a master node is coupled to a Primary Reference Clock (PRC) and a plurality of slave nodes are each arranged to synchronise their internal clock to the PRC using data received on an incoming data link, the method comprising:

30 propagating Synchronisation Status Messages through the network from the master node, with each node through which a message passes incorporating into the message its own identity, thereby generating in each message a path which has been followed by the message;

for each incoming link of each slave node, registering the path or path length of a Synchronisation Status Message received on that link as an attribute for that link; and

- for each slave node, identifying the best incoming signalling link and synchronising on that link after a time delay, indicated by the path of the Synchronisation Status Message received on that link, has elapsed, assuming that no better link is identified in the meantime.

According to a sixth aspect of the present invention there is provided a method of synchronising a node of a telecommunication network in which a master node is coupled to a Primary Reference Clock (PRC) and a plurality of slave nodes are each arranged to synchronise their internal clock to the PRC using data received on incoming data link, the method comprising:

- sending a Synchronisation Status Message Request from the node to be synchronised to neighbouring nodes in the network;
- returning Synchronisation Status Messages from the neighbouring nodes to the requesting node, said messages including a path which has been followed by the message from the master node;
- synchronising said node on the incoming signalling link over which the message having the shortest path was received;
- extending the path of the message having the shortest path to include the identity of the receiving node; and
- propagating the modified synchronisation message to at least certain of the neighbouring nodes.

According to a seventh aspect of the present invention there is provided a node for use in a multi-node telecommunications network, the node comprising:

- means for sending a Synchronisation Status Message Request to neighbouring nodes in the network;
- means for receiving on incoming links to the node, respective Synchronisation Status Messages incorporating the identities of the nodes through which the messages have passed;
- means for registering the paths or path lengths of the Synchronisation Status Messages as attributes for the respective links on which they were received;

means for incorporating into the message having the shortest path length the identity of the node, thereby generating in the message a path which has been followed by the message; and

means for propagating the modified Synchronisation Status Message to at least
5 certain neighbouring nodes using outgoing links.

According to an eighth aspect of the present invention there is provided a method of synchronising a node of a telecommunication network in which a master node is coupled to a Primary Reference Clock (PRC) and a plurality of slave nodes are each
10 arranged to synchronise their internal clock to the PRC using data received on incoming data link, the method comprising:

sending a Synchronisation Status Message Request from the node to be synchronised to neighbouring nodes in the network;

returning Synchronisation Status Messages from the neighbouring nodes to the
15 requesting node, said messages including a path length which has been taken by the message from the master node;

synchronising said node on the incoming signalling link over which the message having the shortest path length was received;

extending the path of the message; and
20 propagating the modified synchronisation message to at least certain of the neighbouring nodes.

Brief Description of the Drawings

25 Figure 1 illustrates schematically a multi-node telecommunication network;
Figure 2 illustrates schematically an alternative multi-node telecommunication network prior to synchronisation;
Figure 3 illustrates the network of Figure 2 following synchronisation;
Figure 4 is a flow diagram illustrating a method of synchronising the nodes of the
30 network of Figures 2 and 3;
Figure 5 illustrates schematically a multi-node telecommunication network prior to synchronisation, and comprising a new node;
Figure 6 illustrates the network of Figure 5 following synchronisation; and

Figure 7 is a flow diagram illustrating a method of synchronising a new node introduced into a multi-node telecommunication network.

Detailed Description of a Preferred Embodiment

5

There is illustrated in Figure 1 a multi-node telecommunication network comprising Nodes A to G. The Nodes are interconnected by data links which may carry user data, signalling data, or a combination of both. In one example, the network of Figure 1 might be a UMTS Terrestrial Radio Access Network (UTRAN), where certain of the nodes (for example Node A) might be Radio Network Controllers (RNCs) whilst others of the nodes (for example Nodes B to G) might be Radio Base Stations (RBSs).

Node A is a so-called "master Node" and is connected to a Primary Reference Clock (PRC). As has already been outlined above, the slave Nodes B to G are able to synchronise with another network Node (and hence with the network as a whole) using data signals received on an incoming data links. The accuracy of the synchronisation will depend to a large extent upon the remoteness of the node which is being synchronised from the master node. An important consideration therefore in choosing which incoming link to synchronise on is the number of inter-node hops which a signal has taken to arrive at the node from the master node.

Upon initialisation of the network of Figure 1, Node A initiates the synchronisation selection process by sending a Synchronisation Status Message (SSM) to each of the nodes to which it is connected (in this case only Node B). The SSM includes a "path" field in which Node A places its own identity together with an indication that Node A is the master node. The SSM is received on a given incoming signalling link by Node B. Node B analyses the SSM and identifies the path. The path is stored as an attribute for the incoming signalling link. Node B then adds its own identity to the path field of the SSM (which becomes {Node A_{PRC}, Node B}), and propagates the modified SSM to Nodes C, D, and E to which it is connected.

The receiving nodes again store the path contained in the received SSM as an attribute for the link on which the message is received. Whilst Nodes D and E are not connected

to any further nodes, Node C is connected to Nodes F and G. Node C therefore adds its identity to the SSM path field (now {Node A_{PRC}, Node B, Node C}) and propagates it to Nodes F and G. Nodes F and G are not connected to any further Nodes and therefore the SSM propagation terminates at these nodes. The path contained in the SSM is stored as an attribute for the incoming links to Nodes F and G.

In the very simple example of Figure 1, each Node has only a single incoming data link on which to synchronise. Figure 2 illustrates a modified network in which an additional link exists between Nodes A and C. In this network, upon initialisation, Node C will receive an SSM from both Nodes A and B. The path contained in the SSM received from Node A will be {Node A_{PRC}} whilst that contained in the SSM received from Node B will be {Node A_{PRC}, Node B}. In the event that the network is fully operational, Node C will select the incoming link from Node A as the link to synchronise on. It does this by comparing the attributes allocated to those links as a result of the respective SSMs. Node C will only choose to synchronise on the incoming link from Node B in the event that the link from Node A fails. Figure 3 illustrates the network of Figure 2 following synchronisation. It will be appreciated that this selection process can be extended to selection from three or more incoming links.

It will also be appreciated that Node C will only propagate to Nodes F and G (at least when the network is fully operational) the SSM which is received from Node A. Only in the event that the link to Node A fails will Node C propagate the SSM received from Node B to Node F and G.

Figure 2 illustrates using dashed lines a so-called "directed loop" which might arise when Node D is connected to Node E and Node E is connected back to Node B. In this situation, an SSM propagated from Node E to Node B will include the path {Node A_{PRC}, Node B, Node D, Node E}. Node B will find that its own identity is contained in the path and hence will detect a directed loop.

30

It will be appreciated that where a Node has several incoming links on which SSMs may be received, it is possible that a node might end up receiving an SSM having a relatively long path before subsequently receiving an SSM having a shorter path. The

5

10

15

20

25

30 This delay provides an opportunity for other SSMs to arrive at the node (and which may have travelled over a shorter path) prior to synchronisation occurring.